

A coincidence detector neural network model of selective attention

Kleanthis Neokleous¹ (kleneokl@cs.ucy.ac.cy)

Maria Koushiou² (maria_koushiou@yahoo.com)

Marios N. Avraamides² (mariosav@ucy.ac.cy)

Christos Schizas¹ (schizas@ucy.ac.cy)

¹Department of Computer Science, ²Department of Psychology
University of Cyprus
P.O Box 20537, Nicosia, Cyprus

Abstract

A computational model of selective attention is implemented to account for findings from an experiment on selective attention that was conducted. The model successfully reproduces the latency data of human participants by relying on the interaction between a bottom-up saliency map and the top-down influences from spatial and semantic goals. The model offers a biologically-plausible way of operationalizing perceptual load and provides insights about the possible brain mechanisms that underlie related empirical findings.

Keywords: selective attention, perceptual load, coincidence detector, neural synchronization.

Introduction

At every moment of conscious life a person experiences a plethora of information that is present in the environment. As processing all available stimulation would represent a daunting task for the perceptual apparatus, a person selects only a subset of information from what is available and discards the rest. The process of focusing mental resources on part of the visual stimulation to allow further processing while ignoring everything else is commonly referred to as selective attention.

Over the years a theoretical debate arose over the locus of selection in the information processing stream. On one hand, early-selection theories of attention (Broadbent, 1958) had posited that selection occurs at an early stage of processing based on the physical characteristics of the stimuli (e.g., intensity, tone). According to these theories, unattended stimuli are discarded from further processing with only some of their physical characteristics being registered. These stimuli are therefore discarded prior to any semantic analysis. On the other hand, late-selection theories (Deutsch & Deutsch, 1963) have placed selection at a later processing stage. These theories have argued that all stimuli enter short-term memory and are thus processed semantically. Those stimuli that are deemed as more important or relevant to the task are then selected and guide response execution.

More recently, Lavie (1995) provided results showing that selection may be early or late depending on task characteristics. In one of Lavie's studies participants were asked to determine whether a presented stimulus was the letter x or the letter z. In one condition, termed the high-load, the letter was flanked by 5 other letters. In another condition, termed the low-load condition, the target was presented accompanied by no flanking letters. In both conditions a distractor letter was presented nearby the target. In a subset of trials the distractor was incompatible to the target designating the alternative response (i.e., if the target was x the distractor was z and vice-versa). In other trials, the distractor was a neutral letter. Results revealed that the interference exerted by the incompatible distractor, evidenced as increased latency for identifying the target, was greater in the low than the high load condition. Lavie accounted for these findings by arguing that in the high-load condition all attentional resources are consumed by the main task leaving no spare resources to process the distractor; that is, the distractor is discarded at an early stage of processing. In contrast, in the low-load condition the task does not exhaust all available resources leaving spare resources to process the distractor. In this case, the distractor is discarded at a later-stage.

Although the Perceptual Load theory offers an appealing account for how selection of information can occur either early or late, it has been recently criticized and challenged. For example, Johnson, McGrath, and McNeil (2002) have shown that an endogenous cue priming the location of the target (i.e., a central arrow presented prior to the target display) eliminates any distractor interference in the low-load condition. As the presentation of a cue does not alter the amount of spare resources, this result is problematic for at least a strong version of the Perceptual Load theory. Furthermore, Elitti, Wallace, and Fox (2005) provided evidence that the critical variable might be the saliency of distractors and not perceptual load per se. In their experiments that manipulated the onsets and offsets of targets and distractors, Elitti et al. showed that interference can be present with high-load provided that the distractor is

made more salient. Finally, Torralbo and Beck (2008) have criticized the Perceptual Load theory on two grounds. First, they argued that the term perceptual load is not clearly defined. Second, they claimed that the concept of exhausted capacity of attentional resources cannot be reconciled easily with what is known about brain mechanisms. Torralbo and Beck (2008) proposed that the neural basis for perceptual load is the extent of competition among stimuli to gain representation in the visual cortex and the strength of a top-down biasing mechanism that is needed to resolve the competition and select a stimulus.

Although Torralbo and Beck (2008) offer a more concrete definition of what constitutes perceptual load, their proposal still seems vague. For example, it is not clear what the nature of the top-down biasing mechanism is and how exactly it operates. In the present study we adopt a computational modelling approach to provide a biologically-plausible account of how perceptual load effects may arise in behavioral experiments. Computational modelling offers an appealing approach towards understanding psychological phenomena as the level of specificity it requires hinders the formulation of vague theories. We have implemented a coincidence detector neural network model of selective attention that accounts for the basic pattern of results obtained in perceptual load experiments, explains how cues may interact with perceptual load, and takes into account the effects of salience. The model is based on recent findings about the neural synchronization between cortical areas during the execution of attentional tasks (Fries, Reynolds, Rorie, & Desimone, 2001) and it also provides for competitive interactions among stimuli at an early stage of processing as suggested by Torralbo and Beck (2008). We compare results from simulations with behavioral data obtained from a perceptual load experiment that involves high and low perceptual load visual searches executed with and without spatial cues. To preview our findings, the model succeeds in accounting for behavioral results, providing thus insights about the possible nature of the neural mechanisms that underlie the perceptual load findings and selective attention in general.

A model of selective attention

Overview

The model is based on evidence that attention is guided by both bottom-up and top-down information (Corbetta & Shulman, 2002). As seen in Figure 1, information in the model is processed at two stages. During the first stage a pre-attentive saliency map is created to represent the low-level features of incoming visual stimuli. At a second stage, the firing rates of neurons in the saliency map are modulated by top-down goals regarding spatial expectations. In addition, a correlation control module assesses the degree of semantic correlation between stimuli and endogenous goals and gradually causes synchronization in neural activity. It should be noted that, in the literature of attention, rate-based

selection and synchronization of neural activity are often proposed as the underlying neural mechanisms of perceptual selection (Buehlmann & Deco, 2008; Niebur, Hsiao, & Johnson, 2002).

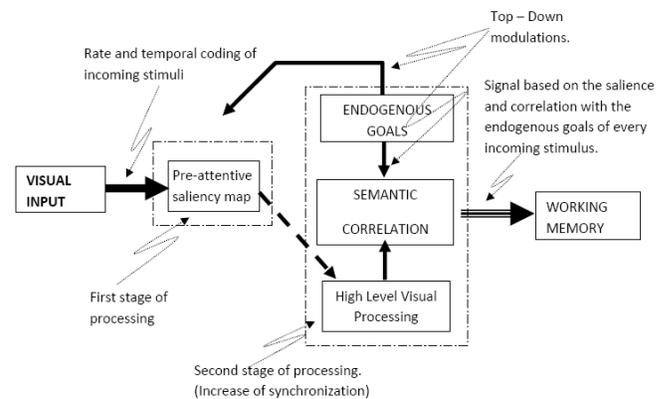


Figure 1. A coincidence detector model of selective attention

Encoding of Stimuli

Encoding in the model is implemented on the basis of the temporal coding hypothesis, that is, the idea that the nervous system uses information about the timing of neural spikes. Synthetic spike trains (i.e., a streams of binary events where the presence or absence of an action potential is represented by 1's and 0's) were produced based on the algorithm proposed by Mikula and Niebur (2008) and Niebur (2007).

Reference spike trains were defined to represent the targets (i.e., letters X and Z). When a visual stimulus entering the receptive field coincided with a target letter, the spike trains of its neurons exhibited a strong correlation with the reference spike train and were therefore more likely to switch their state to that of the reference spike.

Saliency Map

The first stage of processing is responsible for the adjustment of the output firing rate of each neuron in the corresponding receptive field. This is performed based on the salience value of each stimulus in a pre-attentive saliency map that is created on the basis of bottom-up information as suggested by Zhaoping (1999a, 1999b, 2002). According to Zhaoping and Dayan (2006) this saliency map is created in the primary visual cortex (V1) with the receptive field of the most active neuron defining the location that will most likely be selected. This proposal is compatible with findings from the literature showing that the response of a V1 neuron is significantly suppressed by contextual inputs outside but nearby its receptive field (Nothdurft, Gallant, & Van Essen, 1999; Sillito, Grieve, Jones, Cudeiro, & Davis, 1995; Wachtler, Sejnowski, & Albright, 2003). That is, the response of a neuron to its preferred input feature, (e.g., a specific orientation, color, motion direction etc.), is suppressed when similar inputs are present nearby (Zhaoping & Dayan, 2006). Examples of such suppressions have been documented for a variety of features, e.g., iso-orientation suppression (Knierrim,

Gerstner, & van Hemmen, 1998), iso-color suppression (Wachtler et al., 2003), and iso-motion-direction suppression (Jones, Grieve, Wang, & Sillito, 2001). Based on the idea of a pre-attentive saliency map, Zhaoping (2002) implemented a biologically-plausible computational model of V1. The model was successful in simulating performance on a variety of tasks including parallel and conjunctive visual search. Following the proposal of Zhaoping (2002), a bottom-up saliency map in which inhibitory interactions among neurons take place was included in the present model.

Top-down influence and synchronization of neural activity

The firing rates of neurons in the saliency map represent biddings of stimuli for attention. However, as shown in Figure 1 the rates of neurons in the saliency map are modulated by top-down factors. We argue that endogenous goals referring to expectations about the possible location of the upcoming target may modulate the rate of neurons in the saliency map (see also Poghosyan & Ioannides, 2008). This is compatible with recent evidence showing that activity in V1 may appear before the presentation of the stimulus when its position is primed by a cue (Shibata et al., 2008; Silver, Ress, & Heeger, 2007).

In addition to influence from top-down spatial goals, the neural activation of each stimulus is progressively modulated by top-down signals of semantic information. We propose that a correlation control mechanism that includes coincidence detector neurons determines the correlation between semantic goals (e.g., to locate an X or a Z) and the neural activity representing semantic information in processed stimuli (Figure 2). Coincidence detectors represent neurons that fire only if two inputs fire temporally close to each other. This is analogous to the synchronous generation of post synaptic potentials by two input neurons in the brain, which pushes the membrane potential of a target neuron over the threshold required to create an action potential. Based on the degree of correlation between stimuli and semantic goals, an amplification or suppression is exerted on the neural activation of the corresponding stimulus.

The operation of the correlation control module leads to the gradual synchronization of neural activity between neurons representing stimulus information and relevant endogenous goals. The direct connection of top-down attention with synchronization is supported by many recent studies (Gross et al., 2004; Niebur et al., 2002). For example, Saalman, Pigarev, and Vidyasagar (2007) recorded neural activity in the posterior parietal cortex and area V4 of the brain of macaques while they were performing a visual matching task and observed synchronization in the timing of neural activity in the two regions when the animal selectively attended to a location. This has provided evidence that parietal neurons which presumably represent neural activity of the endogenous

goals may selectively increase activity in earlier visual areas.

The neural activation of each stimulus after the first and second stage of processing will provide the input to working memory. A response will be selected based on the stronger activation. The level of activation also determines the latency for selecting the response. To that purpose, working memory consists of recurrent nodes that fire continuously until the activation of one stimulus reaches to a threshold that is needed to select a response.

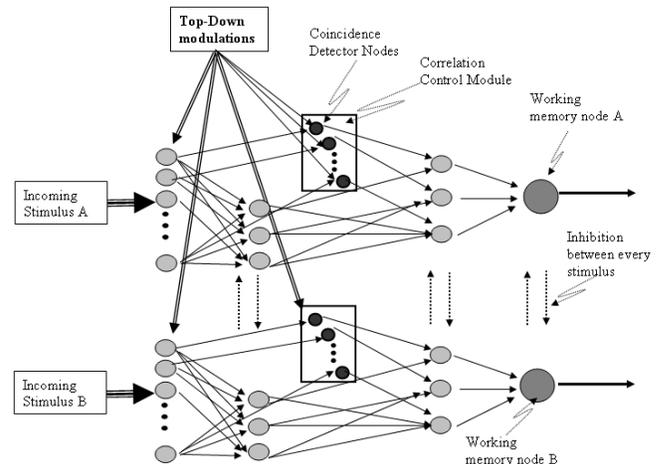


Figure 2. The correlation control module

To implement the model in a perceptual load task, we have conducted a behavioral experiment. The experiment is described next followed by the results from the computational simulations.

Experiment: the effects of perceptual load

A behavioral experiment was carried out in an attempt to (1) replicate the basis pattern of findings obtained with perceptual-load experiments, (2) examine the interaction between cuing and perceptual load, and (3) to generate data that would allow a comparison with the model's output. The experiment conducted was similar to that of Johnson et al., (2002) with the exception that an 80%-valid peripheral cue was used. In their experiment Johnson et al. (2002) have used a 100%-valid central cue.

Method

Participants Twenty-four undergraduate students from the University of Cyprus participated in this experiment in exchange of course credit.

Materials and Design A 2 (perceptual load: high, low) × 3 (cue: no cue, valid, invalid) × 3 (distractor compatibility: neutral, compatible, incompatible) within-subject design was used. The experiment was designed and presented using the E-Prime software package. During testing participants

were seated approximately 50-60cm from a computer screen.

Procedure Participants were asked to perform a visual search task in which 6 letters arranged in a circular array in the center of the screen were presented after a fixation cross. In the high load condition, the search array comprised of the target (X or Z) and 5 letters (M, K, N, H, W) that shared features with the two possible targets (Figure 3). In the low load condition the target appeared among 5 O's (Figure 4).

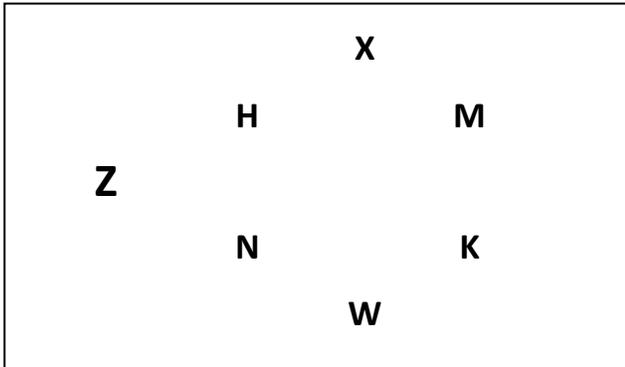


Figure 3: High-load trial with incompatible distractor¹

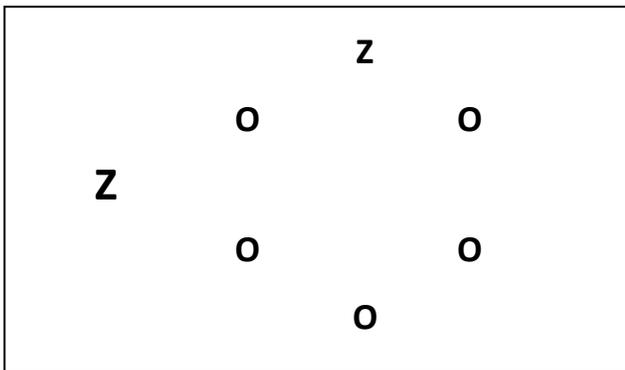


Figure 4: Low-load trial with compatible distractor¹.

In both conditions, a larger distractor letter (X, Z or L) appeared simultaneously on the left or on the right of the circular array and it was compatible (identical to the target), incompatible (calling for the alternative response) or neutral (the letter L) in 1/3 of the trials. Participants were requested to ignore the distractor letter and focus on identifying the target (X or Z) in the 6-letter search array. They were instructed to press “0” for X and “2” for Z as fast as possible using their thumb and their index finger respectively on the numeric keyboard. The search array remained onscreen for 200 ms and was replaced by a mask. Participants were allowed to enter their response for up to 2000 ms after the offset of the search array.

In cue-present trials an asterisk was presented for 150 ms after the offset of the fixation point. The cue was located either in the same location as the target letter (valid cue), or in another position in the circle (invalid cue).

¹ Error bars represent standard errors

Each participant carried out 216 experimental trials, 50% in the low load condition and the other 50% in the high load condition. In each load condition there were 96 valid cue trials (i.e., 80% of the total cue present trials), 96 no cue trials, and 24 invalid cue trials. Reaction time (RT) and accuracy scores were recorded and were used for data analysis.

Results and Discussion

A Repeated Measures ANOVA on median RTs² with load (high vs low), cue (valid cue vs no cue) and compatibility (compatible, incompatible, neutral) as factors was carried out. Invalid cue conditions were excluded from the analysis due to their proportionately small number of experimental trials.

As predicted, the analysis revealed a main effect of load, with participants being faster in the low load than in the high load condition, $F(1,23) = 41.13, p < .001$. Furthermore, latency was shorter for valid than invalid cue trials, $F(1,23) = 58.61, p < .001$. Also, latency was shorter for trials with neutral than either compatible or incompatible distractors, $F(2,46) = 6.23, p < .05$. In addition, a significant load \times cue interaction was obtained, $F(1,23) = 15.69, p < .05$. More importantly though, a significant 3-way interaction (load \times cue \times compatibility) was found, $F(2,46) = 3.43, p < .05$.

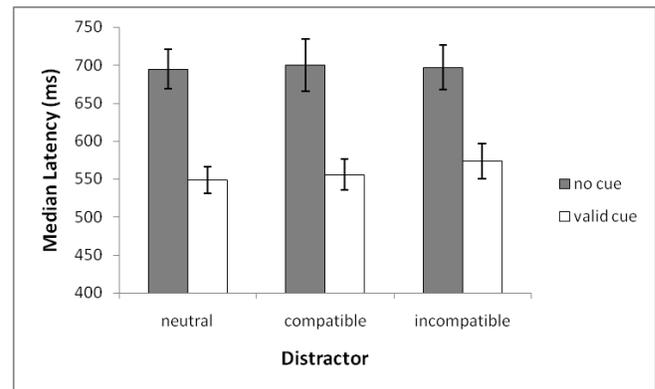


Figure 5: Median latency for high load condition

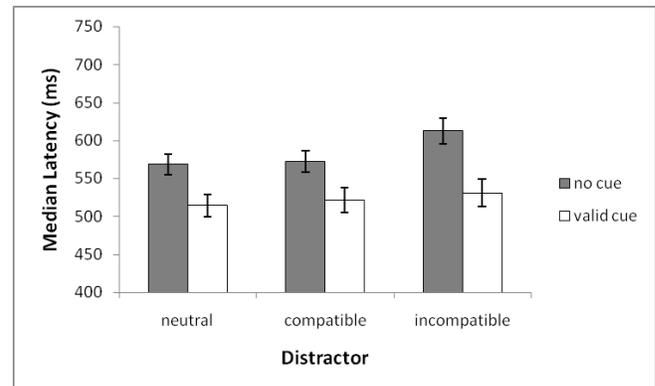


Figure 6: Median Latency for low load condition

² Accuracies were also analyzed but they are not reported here for the sake of brevity

As seen in Figure 5, no difference between the cue conditions was observed in the high load condition, $p = .52$. However, distractor interference in the low load condition was significantly reduced when a valid cue was provided (Figure 6). Indeed, a paired sample t-test revealed a significant difference between incompatible and neutral trials in the low load condition, $t(23) = 4.04, p < .001$.

The pattern of results obtained in the experiment replicated the typical pattern of findings of perceptual-load studies (Lavie, 1995). Furthermore, it showed that cueing interacts with perceptual load. Johnson et al., (2002) have shown that a 100% valid central cue diminished interference from incompatible distractors in the low-load condition, equating in fact the extent of interference in the low and high load conditions. The present experiment provided evidence that an 80% valid peripheral cue was similarly effective. The present results are incompatible with at least a strong version of the perceptual-load theory as they show that even in low-load conditions, in which spare resources exist to process a distractor, interference may be absent.

Computational simulations

The coincidence detector model of selective attention was used to simulate performance in all conditions included in the experiment. The basic pattern of the perceptual load findings is simulated by means of the inhibitory interactions among neurons representing stimuli in the saliency map and the top-down spatial and semantic influences.

In the high load condition the encoding mechanisms adopted in model yielded similar neural activity for each stimulus in the circular array (i.e., target and flankers) while in the low load condition the neuron representing the target had enhanced activity relative to its flankers. This is because target and flankers share more features in the high load condition than in the low load condition; this is compatible with what is known about iso-feature suppression (see Zhaoping, 2002). There is no effect of an incompatible distractor in the high load condition because, although it has high activation from its correlation with semantic goals, the distractor also shares features with all stimuli on the screen. As a result, it competes with all other items and is therefore prone to inhibition from many sources. In contrast, the effect of an incompatible distractor in the low load condition is high. This is because (1) the distractor has high activation due to its semantic correlation with the endogenous goals, and (2) it only competes with the target as it shares no features with the flankers.

The effects of a spatial cue were modeled by adding a spatial goal (i.e., “look first for the item that will appear in this location”). In the simulations, the item that appeared in the cued location received a boost in its activity due to its correlation with the spatial goal. All other items received inhibition.

Figures 7 and 8 present the predictions of the model after 100 runs in each condition.

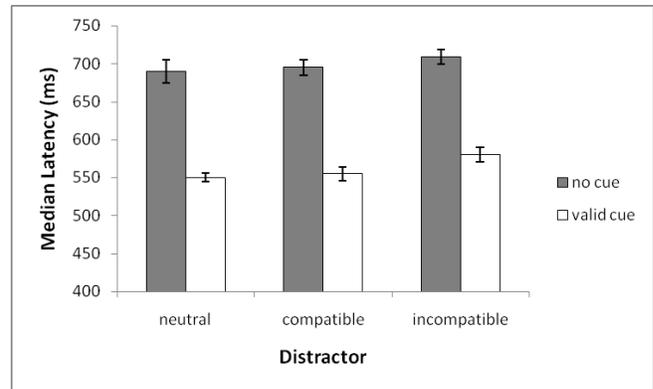


Figure 7: Simulation results for high load condition

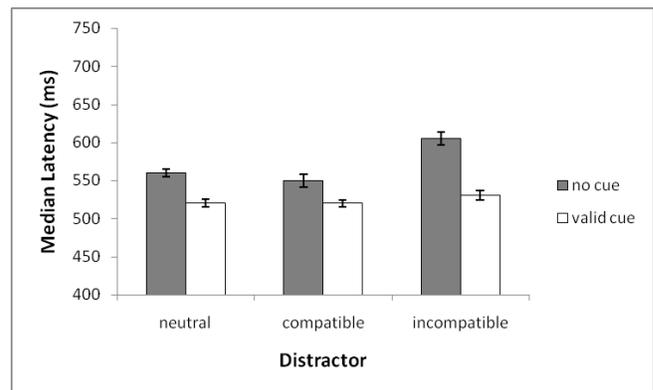


Figure 8: Simulation results for low load condition

Conclusion

The computational model of selective attention implemented here was able to account for the findings from the behavioral experiment on perceptual load which produced not only the basic pattern of findings reported by other studies (Lavie, 1995) but also the findings on the interaction of spatial cueing with perceptual load (Johnson et al., 2002). Thus, the model offers an explicit hypothesis about the possible neural mechanisms that give rise to these findings without relying on a vague use of the term perceptual load. To that respect, the model is compatible with the conjecture of Torralbo and Beck (2008) that perceptual load effects are a product of the competitive interactions among stimuli for neural representation. In addition, the model suggests that the correlation between input and endogenous goals (spatial and semantic) may represent the biasing mechanism proposed by Torralbo and Beck (2008) as the means for resolving the competition.

Importantly the model was designed based on what is currently known about the underlying neural mechanisms of selective attention. The presence of a saliency map, the modulation of firing rate of neurons in V1, and the synchronization of activity in populations of neurons during selective attention are all supported by research in psychology and cognitive neuroscience. Clearly, further research is needed to determine which of these mechanisms provide indeed an index of selective attention. Also, more research is needed to evaluate the validity of the presented

model. To that respect, it is important for future work to examine whether the model is capable, without any tuning, to account for other findings from the attention literature.

Acknowledgments

This research is supported by a research grant from the Cyprus Research Promotion Foundation. We thank Vahe Poghosyan for his insightful comments and suggestions.

References

- Broadbent, D. E. (1958). *Perception and communication*. Oxford: Pergamon.
- Buehlmann, A., & Deco, G. (2008). The neuronal basis of attention: rate versus synchronization modulation. *The Journal of Neuroscience*, *28*, 7679-7686.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*, 201-215.
- Deutsch, J. A., & Deutsch, D. (1963). Attention: Some theoretical considerations. *Psychological Review*, *1963*, 283-321.
- Elitti, S., Wallace, D., & Fox, E. (2005). Selective target processing: perceptual load or distractor salience. *Perception & Psychophysics*, *67*, 876-885.
- Fries, A. K., Reynolds, J. H., Rorie, A. E., & Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, *291*, 1560-1563.
- Gross, J., Schmitz, F., Schnitzler, I., Kessler, K., Shapiro, K., Hommel, B., et al. (2004). Modulation of long-range neural synchrony reflects temporal limitations of visual attention in humans. *Proceedings of the National Academy of Sciences of United States of America*, *101*, 13050-13055.
- Johnson, D., McGrath, A., & McNeil, C. (2002). Cuing interacts with perceptual load in visual search. *Psychological Science*, *13*, 284-287.
- Jones, H. E., Grieve, K. L., Wang, W., & Sillito, M. A. (2001). Surround suppression in primate V1. *The Journal of Neurophysiology*, *86*, 2011-2028.
- Knierim, J. J., Gerstner, W., & van Hemmen, J. (1998). Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *Journal of Neurophysiology*, *67*, 961-980.
- Lavie, N. (1995). Perceptual load as a necessary condition for selective attention. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 451-468.
- Mikula, S., & Niebur, E. (2008). Exact Solutions for Rate and Synchrony in Recurrent Networks of Coincidence Detectors. *Neural Computation*, *20*, 2637-2661.
- Niebur, E. (2007). Generation of synthetic spike trains with defined pairwise correlations. *Neural Computation*, *19*, 1720-1738.
- Niebur, E., Hsiao, S. S., & Johnson, K. O. (2002). Synchrony: a neuronal mechanism for attentional selection? *Current Opinion in Neurobiology*, *12*, 190-194.
- Nothdurft, H. C., Gallant, J. L., & Van Essen, D. C. (1999). Response modulation by texture surround in primate area V1: correlates of "popout" under anesthesia. *Visual Neuroscience*, *16*, 15-34.
- Poghosyan V, & Ioannides, A. A. (2008). Attention modulates earliest responses in the primary auditory and visual cortices. *Neuron*, *58*, 802-13.
- Saalman, Y. B., Pigarev, I. N., & Vidyasagar, T. R. (2007). Neural Mechanisms of Visual Attention: How Top-Down Feedback Highlights Relevant Locations. *Science*, *316*, 1612-1615.
- Shibata, K., Yamagishi, N., Goda, N., Yoshioka, T., Yamashita, O., Sato, M. A., et al. (2008). The effects of feature attention on prestimulus cortical activity in the human visual system. *Cerebral Cortex*, *18*, 1664-1675.
- Sillito, M. A., Grieve, K. L., Jones, H. E., Cudeiro, J., & Davis, J. (1995). Visual cortical mechanisms detecting focal orientation discontinuities. *Nature* *378*, 492-496.
- Silver, M. A., Ress, D., & Heeger, D. J. (2007). Neural correlates of sustained spatial attention in human early visual cortex. *The Journal of Neurophysiology*, *97*, 229-237.
- Torralbo, A., & Beck, D. M. (2008). Perceptual-Load-induced selection as a result of local competitive interactions in visual cortex. *Psychological Science*, *19*, 1045-1050.
- Wachtler, T., Sejnowski, T. J., & Albright, T. D. (2003). Representation of color stimuli in awake macaque primary visual cortex. *Neuron*, *37*, 681-691.
- Zhaoping, L. (1999a). Contextual influences in V1 as a basis for pop out and asymmetry in visual search. *Proceedings of the National Academy of Sciences of United States of America*, *96*, 10530-10535.
- Zhaoping, L. (1999b). Visual segmentation by contextual influences via intracortical interactions in primary visual cortex. *Network: Computation and Neural Systems*, *10*, 187-212.
- Zhaoping, L. (2002). A saliency map in primary visual cortex. *Trends in Cognitive Sciences*, *9-15*.
- Zhaoping, L., & Dayan, P. (2006). Pre-attentive visual selection. *Neural Networks*, *19*, 1437-1439.